arXiv:1405.4356v2 [cs.DC] 19 Jun 2014

Lessons from the Congested Clique Applied to MapReduce ⋆

James W. Hegeman and Sriram V. Pemmaraju

Department of Computer Science
The University of Iowa
Iowa City, Iowa 52242-1419, USA

Abstract. The main results of this paper are (I) a simulation algorithm which, under quite general constraints, transforms algorithms running on the Congested Clique into algorithms running in the MapReduce model, and (II) a distributed $O(\Delta)$-coloring algorithm running on the Congested Clique which has an expected running time of $O(1)$ rounds, if $\Delta \geq \Theta(\log^4 n)$; and $O(\log\log\log n)$ rounds otherwise. Applying the simulation theorem to the Congested Clique $O(\Delta)$-coloring algorithm yields an $O(1)$-round $O(\Delta)$-coloring algorithm in the MapReduce model.

Our simulation algorithm illustrates a natural correspondence between per-node bandwidth in the Congested Clique model and memory per machine in the MapReduce model. In the Congested Clique (and more generally, any network in the $\mathcal{CONGEST}$ model), the major impediment to constructing fast algorithms is the $O(\log n)$ restriction on message sizes. Similarly, in the MapReduce model, the combined restrictions on memory per machine and total system memory have a dominant effect on algorithm design. In showing a fairly general simulation algorithm, we highlight the similarities and differences between these models.

1 Introduction

The $\mathcal{CONGEST}$ model of distributed computation is a synchronous, message-passing model in which the amount of information that a node can transmit along an incident edge in one round is restricted to $O(\log n)$ bits [15]. As the name suggests, the $\mathcal{CONGEST}$ model focuses on *congestion* as an obstacle to distributed computation. Recently, a fair amount of research activity has focused on the design of distributed algorithms in the $\mathcal{CONGEST}$ model assuming that the underlying communication network is a *clique* [2,5,12,14]. Working with such a *Congested Clique* model completely removes from the picture obstacles that might be due to nodes having to acquire information from distant nodes (since any two nodes are neighbors), thus allowing us to focus on the problem of congestion alone. Making this setting intriguing is also the fact that no non-trivial lower bounds for computation on a Congested Clique have been proved. In fact, in a recent paper, Lenzen [12] showed how to do load-balancing deterministically so as to route up to n^2 messages (each of size $O(\log n)$) in $O(1)$ rounds in the Congested Clique setting, provided each node is the source of at most n messages and the sink for at most n messages. Thus a large volume of information can be moved around the network very quickly and any lower-bound approach in the Congested Clique setting will have to work around Lenzen's routing-protocol result. While Lotker et al. [13] mention overlay networks as a possible practical application of distributed computation on a Congested Clique, as of now, research on this model is largely driven by a theoretical interest in exploring the limits imposed by congestion.

MapReduce [4] is a tremendously popular parallel-programming framework that has become the tool of choice for large-scale data analytics at many companies such as Amazon, Facebook, Google, Yahoo!, etc., as well as at many universities. While the actual time-efficiency of a particular MapReduce-like implementation will depend on many low-level technical details, Karloff et al. [9] have attempted to formalize key constraints of this framework to propose a *MapReduce model* and an associated MapReduce complexity class ($\mathcal{MRC}$). Informally speaking, a problem belongs to $\mathcal{MRC}$ if it can be solved in the MapReduce framework using: (i) a number of machines that is substantially sublinear in the input size, i.e., $O(n^{1-\epsilon})$ for constant $\epsilon > 0$, (ii) memory per machine that is substantially sublinear in the input size, (iii) $O(\text{poly}(\log n))$ number of map-shuffle-reduce rounds, and (iv) polynomial-time local computation at each machine in each round. Specifically, a problem is said to be in $\mathcal{MRC}^i$ if it can be solved in $O(\log^i n)$ map-shuffle-reduce rounds, while maintaining the other constraints mentioned above. Karloff et al. [9] show that *minimum spanning tree* (MST) is in $\mathcal{MRC}^0$ (i.e., MST requires $O(1)$ map-shuffle-reduce rounds) on non-sparse instances. Following up on this, Lattanzi et al. [11] show that other problems such as *maximal matching* (with which the distributed computing community is very familiar) are also in $\mathcal{MRC}^0$ (again, for non-sparse instances). We give a more-detailed description of the MapReduce model in Section 1.1.

⋆ This work is supported in part by National Science Foundation grant CCF 1318166. E-mail: [james-hegeman,sriram-pemmaraju]@uiowa.edu.

The volume of communication that occurs in a Shuffle step can be quite substantial and provides a strong incentive to design algorithms in the MapReduce framework that use very few map-shuffle-reduce steps. As motivation for their approach (which they call *filtering*) to designing MapReduce algorithms, Lattanzi et al. [11] mention that past attempts to "shoehorn message-passing style algorithms into the framework" have led to inefficient algorithms. While this may be true for distributed message-passing algorithms in general, we show in this paper that algorithms designed in the Congested Clique model provide many lessons on how to design algorithms in the MapReduce model. We illustrate this by first designing an expected-$O(1)$-round algorithm for computing a $O(\Delta)$-coloring for a given n-node graph with maximum degree $\Delta \geq \log^4 n$ in the Congested Clique model. We then simulate this algorithm in the MapReduce model and obtain a corresponding algorithm that uses a constant number of map-shuffle-reduce rounds to compute an $O(\Delta)$-coloring of the given graph. While both of these results are new, what we wish to emphasize in this paper is the *simulation* of Congested Clique algorithms in the MapReduce model. Our simulation can also be used to obtain efficient MapReduce-model algorithms for other problems such as 2-*ruling sets* [2] for which an expected-$O(\log\log n)$-round algorithm on a Congested Clique was recently developed.

1.1 Models

The Congested Clique Model. The Congested Clique is a variation on the more general $\mathcal{CONGEST}$ model. The underlying communication network is a size-n clique, i.e., every pair of nodes can directly communicate with each other. Computation proceeds in synchronous rounds and in each round a node (i) receives all messages sent to it in the previous round; (ii) performs unlimited local computation; and then (iii) sends a, possibly distinct, message of size $O(\log n)$ to each other node in the network. We assume that nodes have distinct IDs that can each be represented in $O(\log n)$ bits. We call this the *Congested Clique* model.

Our focus in this paper is graph problems and we assume that the input is a graph G that is a spanning subgraph of the communication network. Initially, each node in the network knows who its neighbors are in G. Thus knowledge of G is distributed among the nodes of the network, with each node having a particular local view of G. Note that G can be quite dense (e.g., have $\Omega(n^2)$ edges) and therefore any reasonably fast algorithm for the problem will have to be "truly" distributed in the sense that it cannot simply rely on shipping off the problem description to a single node for local computation.

The MapReduce Model. Our description of the MapReduce model borrows heavily from the work of Karloff et al. [9] and Lattanzi et al. [11]. Introduced by Karloff et al. [9], the MapReduce model is an abstraction of the popular MapReduce framework [4] implemented at Google and also in the popular Hadoop open-source project by Apache.

The basic unit of information in the MapReduce model is a $(key, value)$-pair. At a high level, computation in this model can be viewed as the application of a sequence of functions, each taking as input a collection of $(key, value)$-pairs and producing as output a new collection of $(key, value)$-pairs. MapReduce computation proceeds in rounds, with each round composed of a map phase, followed by a shuffle phase, followed by a reduce phase. In the map phase, $(key, value)$ pairs are processed individually and the output of this phases is a collection of $(key, value)$-pairs. In the shuffle phase, these $(key, value)$-pairs are "routed" so that all $(key, value)$-pairs with the same key end up together. In the last phase, namely the reduce phase, each key and all associated values are processed together. We next describe each of the three phases in more detail.

- The computation in the Map phase of round i is performed by a collection of *mappers*, one per $(key, value)$ pair. In other words, each mapper takes a $(key, value)$ pair and outputs a collection of $(key, value)$ pairs. Since each mapper works on an individual $(key, value)$ pair and the computation is entirely "stateless" (i.e., not dependent on any stored information from previous computation), the mappers can be arbitrarily distributed among machines. In the MapReduce model, keys and values are restricted to the word size of the system, which is $\Theta(\log n)$. Because of this restriction, a mapper takes as input only a constant number of words.
- In the *Shuffle* phase of round i, which runs concurrently with the Map phase (as possible), key-value pairs emitted by the mappers are moved from the machine that produced them to the machine which will run the reducer for which they are destined; i.e., a key-value pair (k, v) emitted by a mapper is physically moved to the machine which will run the reducer responsible for key k in round i. The Shuffle phase is implemented entirely by the underlying MapReduce framework and we generally ignore the Shuffle phase and treat data movement from one machine to another as a part of the Map phase.
- In the *Reduce* phase of round i, reducers operate on the collected key-value pairs sent to them; a reducer is a function taking as input a pair $(k, \{v_{k,j}\}_j)$, where the first element is a key k and the second is a multiset of values $\{v_{k,j}\}_j$ which comprises all of the values contained in key-value pairs emitted by mappers during round i and

having key k. Reducers emit a multiset of key-value pairs $\{(k, v_{k,l})\}_l$, where the key k in each pair is the same as the key k of the input.

For our purposes, the concepts of a machine and a reducer are interchangeable, because reducers are allowed to be "as large" as a single machine on which they compute.

The MapReduce model of Karloff et al. [9] tries to make explicit three key resource constraints on the MapReduce system. Suppose that the problem input has size n (note that this is *not* referring to the input size of a particular reducer or mapper). We assume, as do Karloff et al. [9] and Lattanzi et al. [11], that memory is measured in $O(\log n)$-bit-sized words.

1. Key-sizes and value-sizes are restricted to a $\Theta(1)$ multiple of the word size of the system. Because of this restriction, a mapper takes as input only a constant number of words.
2. Both mappers and reducers are restricted to using space consisting of $O(n^{1-\epsilon})$ words of memory, and time which is polynomial in n.
3. The number of machines, or equivalently, the number of reducers, is restricted to $O(n^{1-\epsilon})$.

Given these constraints, the goal is to design MapReduce algorithms that run in very few – preferably constant – number of rounds. For further details on the justifications for these constraints, see [9].

Since our focus is graph algorithms, we can restate the above constraints more specifically in terms of graph size. Suppose that an n-node graph $G = (V, E)$ is the input. Following Lattanzi et al. [11], we assume that each machine in the MapReduce system has memory $\eta = n^{1+\epsilon}$ for $\epsilon \geq 0$. Since $n^{1+\epsilon}$ needs to be "substantially" sublinear in the input size, we assume that the number of edges m of G is $\Omega(n^{1+c})$ for $c > \epsilon$. Thus the MapReduce results in this paper are for non-sparse graphs.

1.2 Contributions

The main contribution of this paper is to show that fast algorithms in the Congested Clique model can be translated via a simulation theorem into fast algorithms in the MapReduce framework. As a case study, we design a fast graph-coloring algorithm running in the Congested Clique model and then apply the simulation theorem to this algorithm and obtain a fast MapReduce algorithm. Specifically, given an n-node graph G with maximum degree $\Delta \geq \log^4 n$, we show how to compute an $O(\Delta)$-coloring of G in expected $O(1)$ rounds in the Congested Clique model. We also present an algorithm for small Δ; for $\Delta < \log^4 n$ we present an algorithm that computes a $\Delta + 1$ coloring in $O(\log \log \log n)$ rounds with high probability on a Congested Clique. The implication of this result to the MapReduce model (via the simulation theorem) is that for any n-node graph with $\Omega(n^{1+c})$ edges, for constant $c > 0$, there is a MapReduce algorithm that runs in $O(1)$ map-shuffle-reduce rounds using $n^{1+\epsilon}$ memory per machine, for $0 \leq \epsilon < c$ and $n^{c-\epsilon}$ machines. Note that the even when using n memory per machine and n^c machines the algorithm still takes $O(1)$ rounds. This is in contrast to examples in Lattanzi et al. [11] such as maximal matching which require $O(\log n)$ rounds if the memory per machine is n.

The coloring algorithms in both models are new and faster than any known in the respective models, as far as we know. However, the bigger point of this paper is the connection between models that are studied in somewhat different communities.

1.3 Related Work

The earliest interesting algorithm in the Congested Clique model is an $O(\log \log n)$-round deterministic algorithm to compute a minimum spanning tree (MST), due to Lotker et al. [13]. Gehweiler et al. [7] presented a random $O(1)$-round algorithm in the Congested Clique model that produced a constant-factor approximation algorithm for the *uniform* metric facility location problem. Berns et al. [2,3] considered the more-general non-uniform metric facility location in the Congested Clique model and presented a constant-factor approximation running in expected $O(\log \log n)$ rounds. Berns et al. reduce the metric facility location problem to the problem of computing a 2-ruling set of a spanning subgraph of the underlying communication network and show how to solve this in $O(\log \log n)$ rounds in expectation. In 2013, Lenzen presented a routing protocol to solve a problem called an *Information Distribution Task* [12]. The setup for this problem is that each node $i \in V$ is given a set of $n' \leq n$ messages, each of size $O(\log n)$, $\{m_i^1, m_i^2, \ldots, m_i^{n'}\}$, with destinations $d(m_i^j) \in V$, $j \in \{1, 2, \ldots, n'\}$. Messages are globally lexicographically ordered by their source i, destination $d(m_i^j)$, and j. Each node is also the destination of at most n messages. Lenzen's routing protocol solves the Information Distribution Task in $O(1)$ rounds.

Our main sources of reference on the MapReduce model and for graph algorithms in this model are the work of Karloff et al. [9] and Lattanzi et al. [11] respectively. Besides these, the work of Ene et al. [6] on algorithms for clustering in MapReduce model and the work of Kumar et al. [10] on greedy algorithms in the MapReduce model are relevant.

2 Coloring on the Congested Clique

In this section we present an algorithm, running in the Congested Clique model, that takes an n-node graph G with maximum degree Δ and computes an $O(\Delta)$-coloring in expected $O(\log\log\log n)$ rounds. In fact, for high-degree graphs, i.e., when $\Delta \geq \log^4 n$, our algorithm computes an $O(\Delta)$-coloring in $O(1)$ rounds. This algorithm, which we call Algorithm HIGHDEGCOL, is the main contribution of this section. For graphs with maximum degree less than $\log^4 n$ we appeal to an already-known coloring algorithm that computes a $(\Delta+1)$ coloring in $O(\log\Delta)$ rounds and then modify its implementation so that it runs in $O(\log\log\log n)$ rounds on a Congested Clique.

We first give an overview of Algorithm HIGHDEGCOL. The reader is advised to follow the pseudocode given in Algorithm 1 as they read the following. The algorithm repeatedly performs a simple random trial until a favorable event occurs. Each trial is independent of previous trials. The key step of Algorithm HIGHDEGCOL is that each node picks a *color group* k from the set $\{1, 2, \ldots, \lceil \Delta/\log n \rceil\}$ independently and uniformly at random (Step 4). We show (in Lemma 1) that the expected number of edges in the graph G_k induced by nodes in color group k is at most $O(\frac{n\log^2 n}{\Delta})$. Of course, some of the color groups may induce far more edges and so we define a *good* color group as one that has at most n edges. The measure of whether the random trial has succeeded is the number of good color groups. If most of the color groups are good, i.e., if at most $2\log n$ color groups are not good then the random trial has succeeded and we break out of the loop. We then transmit each graph induced by a good color group to a distinct node in constant rounds using Lenzen's routing scheme [12] (Step 11). Note that this is possible because every good color group induces a graph that requires $O(n)$ words of information to completely describe. Every node that receives a graph induced by a good color group locally computes a proper coloring of the graph using one more color than the maximum degree of the graph it receives (Step 12). Furthermore, every such coloring in an iteration employs a distinct palette of colors. Since there are very few color groups that are not good, we are able to show that the residual graph induced by nodes not in good color groups has $O(n)$ edges. As a result, the residual graph can be communicated in its entirety to a single node for local processing. This completes the coloring of all nodes in the graph.

We now analyze Algorithm HIGHDEGCOL and show that (i) it terminates in expected-$O(1)$ rounds and (ii) it uses $O(\Delta)$ colors. Subsequently, we discuss an $O(\log\log\log n)$ algorithm to deal with the small Δ case.

Lemma 1. *For each k, the expected number of edges in G_k is $\frac{n\log^2 n}{2\Delta}$.*

Proof: Consider edge $\{u, v\}$ in G. The probability that both u and v choose color group k is at most $\frac{\log n}{\Delta} \cdot \frac{\log n}{\Delta} = \frac{\log^2 n}{\Delta^2}$. Since G has at most $\frac{1}{2}\Delta \cdot n$ edges, the expected number of edges in G_k is at most $\frac{n\log^2 n}{2\Delta}$. □

Lemma 2. *The expected number of color-group graphs G_k having more than n edges is at most $\log n$.*

Proof: By Lemma 1 and Markov's inequality, the probability that color group k has more than n edges is at most $\frac{n\log^2 n}{2\Delta\cdot n} = \frac{\log^2 n}{2\Delta}$. Since there are $\lceil \Delta/\log n \rceil$ groups, the expected number of G_k having more than n edges is bounded above by $2\frac{\Delta}{\log n} \cdot \frac{\log^2 n}{2\Delta} = \log n$. □

Lemma 3. *With high probability, every color group has $\frac{5n\log n}{\Delta}$ nodes.*

Proof: The number of color groups is $\lceil \Delta/\log n \rceil$. Thus, for any k, the expected number of nodes in G_k, denoted $|V(G_k)|$, is at most $n \cdot \frac{\log n}{\Delta}$. An application of a Chernoff bound then gives, for each k,

$$\mathbf{P}\left(|V(G_k)| > 5n \cdot \frac{\log n}{\Delta}\right) \leq 2^{-5n\cdot\frac{\log n}{\Delta}} < 2^{-5\log n} = \frac{1}{n^5}$$

Taking the union over all k completes the proof. □

Lemma 4. *With high probability, no node u in G has more than $5\log n$ neighbors in any color group.*

Algorithm 1 HIGHDEGCOL

Input: An n-node graph $G = (V, E)$, of maximum degree Δ
Output: A proper node-coloring of G using $O(\Delta)$ colors

1. Each node u in G computes and broadcasts its degree to every other node v in G.
2. **If** $\Delta \leq \log^4 n$ **then** use Algorithm LOWDEGCOL instead.
3. **while** $true$ **do**
4. Each node u chooses a *color group* k from the set $\{1, 2, \ldots, \lceil \Delta / \log n \rceil\}$ independently and uniformly at random.
5. Let G_k be the subgraph of G induced by nodes of color group k.
6. Each node u sends its choice of color group to all neighbors in G.
7. Each node u computes its degree within its own color-group graph G_{k_u} and sends its color group and degree within color group to node 1.
8. Node 1, knowing the partition of G into color groups and also knowing the degree of every node u ($u \in G_k$) within the induced subgraph G_k, can compute the number of edges in G_k for each k. Thus node 1 can determine which color-group graphs G_k are *good*, i.e., have at most n edges.
9. If at most $2 \log n$ color-group graphs are not good, node 1 broadcasts a "break" message to all nodes causing them to **break** out of loop;

endwhile

10. Node 1 informs every node u in a good group of the fact that u's color group is good
11. Using Lenzen's routing protocol, distribute all information about all good color-group graphs G_k to distinct nodes of G.
12. For each good G_k, the recipient of G_k computes a coloring of G_k using $\Delta(G_k) + 1$ colors. The color palette used for each G_k is distinct.
13. The residual graph $\overline{G}$ of uncolored nodes has size $O(n)$ with high probability, and can thus be transmitted to a single node (for local proper coloring) in $O(1)$ rounds.
14. Each node that locally colors a subgraph informs each node in the subgraph the color it has been assigned.

Proof: Node u has maximum degree Δ, so for any k, the expected number of neighbors of u which choose color group k is bounded above by $\log n$. Therefore, applying a Chernoff bound gives

$$\mathbf{P}\left(|N(u) \cap G_k| > 5 \log n\right) \leq 2^{-5 \log n} = \frac{1}{n^5}$$

Taking the union over all k and u shows that, with probability at least $1 - \frac{1}{n^3}$, the assertion of the lemma holds. □

Lemma 5. *The residual graph $\overline{G}$, induced by groups that are good, has $O(n)$ edges, with high probability.*

Proof: The residual graph $\overline{G}$ is a graph induced by at most $2 \log n$ color groups, since the algorithm is designed to terminate only when it has performed a trial resulting in at most $2 \log n$ groups that are not good. With high probability, no node u in $\overline{G}$ has more than $5 \log n$ neighbors in any of the (at most) $2 \log n$ color groups that make up $\overline{G}$, so therefore with high probability no node u has degree greater than $10 \log^2 n$ in $\overline{G}$. Since $\overline{G}$ has at most $(2 \log n) \cdot \frac{5n \log n}{\Delta}$ nodes with high probability, it follows that the number of edges in $\overline{G}$ is at most

$$(2 \log n) \cdot \frac{5n \log n}{\Delta} \cdot 10 \log^2 n = \frac{100 n \log^4 n}{\Delta}$$

which is $O(n)$ when $\Delta \geq \log^4 n$. □

Lemma 6. *Algorithm* HIGHDEGCOL *runs in a constant number of rounds, in expectation.*

Proof: By Lemma 2 and Markov's inequality, the expected number of color-group partitioning attempts required before the number of "bad" color groups (i.e., color groups whose induced graphs G_k contain more than n edges) is less than or equal to $2 \log n$ is two. It is easy to verify that each iteration of the **while**-true loop requires $O(1)$ rounds of communication.

When $\Delta \geq \log^4 n$, the residual graph $\overline{G}$ is of size $O(n)$ with high probability, and can thus be communicated in its entirety to a single node in $O(1)$ rounds. That single node can then color $\overline{G}$ deterministically using $\Delta + 1$ colors and then inform every node of $\overline{G}$ of its determined color in one further round. □

Lemma 7. *Algorithm* HIGHDEGCOL *uses $O(\Delta)$ colors.*

Proof: A palette of size $O(\log n)$ colors suffices for each good color group because we showed in Lemma 4 that with high probability the maximum degree in any color group is $5 \log n$. Since there are a total of $\lceil \Delta / \log n \rceil$ color groups and we use a distinct palette of size $O(\log n)$ for each good color group, we use a total of $O(\Delta)$ colors for the good color groups. The residual graph induced by not-good color groups is colored in the last step and it requires an additional $O(\Delta)$ colors. □

2.1 Coloring low-degree graphs

Now we describe an algorithm that we call LOWDEGCOL that, given an n-node graph G with maximum degree $\Delta < \log^4 n$, computes a proper $(\Delta + 1)$-coloring with high probability in $O(\log \log \log n)$ rounds in the Congested Clique model. The algorithm has two stages. The first stage of the algorithm is based on the simple, natural, randomized coloring algorithm first analyzed by Johannson [8] and more recently by Barenboim et al. [1]. Each node u starts with a color palette $C_u = \{1, 2, \ldots, \Delta + 1\}$. In each iteration, each as-yet uncolored node u makes a tentative color choice $c(u) \in C_u$ by picking a color from C_u independently and uniformly at random. If no node in u's neighborhood picks color $c(u)$ then u colors itself $c(u)$ and $c(u)$ is deleted from the palettes of all neighbors of u. Otherwise, u remains uncolored and participates in the next iteration of the algorithm. We call one such iteration RANDCOLSTEP. Barenboim et al. [1] show (as part of the proof of Theorem 5.1) that if we executed $O(\log \Delta)$ iterations of RANDCOLSTEP, then with high probability the nodes that remain uncolored induce connected components of size $O(\text{poly}(\log n))$. Since we are evaluating a situation in which $\Delta < \log^4 n$, this translates to using $O(\log \log n)$ iterations of RANDCOLSTEP to reach a state with small connected components. Now notice that this algorithm uses only the edges of G – the graph being colored – for communication. By utilizing the entire bandwidth of the underlying clique communication network, it is possible to speed up this algorithm significantly and get it to complete in $O(\log \log \log n)$ rounds. The trick to doing this is to rapidly gather, at each node u, all information needed by node u to execute the algorithm locally. We make this precise further below.

Once we execute $O(\log \log \log n)$ iterations of RANDCOLSTEP and all connected components induced by as-yet uncolored nodes become polylogarithmic in size, then Stage 2 of the algorithm begins. In this stage, first each connected component is gathered at a node; we show how to accomplish this in $O(\log \log \log n)$ rounds by appealing to the deterministic MST algorithm on a Congested Clique due to Lotker et al. [13]. Then each connected component of uncolored nodes is shipped off to a distinct node and is locally (and independently) colored using $\Delta + 1$ colors.

We start by developing Stage 1 first. Suppose that for some constants c_1, c_2, c_3, $T < c_1 \log \log n$ iterations of RANDCOLSTEP are needed before all connected components induced by uncolored nodes have size at most $c_2 \cdot \log^{c_3} n$ with probability at least $1 - 1/n$. Let G_L denote a labeled version of graph G in which each node u is labeled $(\mathtt{ID}_u, \mathtt{RS}_u)$, where $\mathtt{ID}_u$ is the $O(\log n)$-bit ID of node u and $\mathtt{RS}_u$ is a random bit string of length $T \cdot \lceil \log \Delta \rceil$. For integer $k \geq 0$ and node $u \in V$, let $B(u, k)$ denote the set of all nodes within k hops of u in G. The following lemma shows that it is quite helpful if each node u knew $G_L[B(u, T)]$, the subgraph of the labeled graph G_L induced by nodes in $B(u, T)$.

Lemma 8. *Suppose that each node $u \in V$ knows $G_L[B(u, T)]$. Then each node u can locally compute a color $c(u) \in \{\perp\} \cup \{1, 2, \ldots, \Delta + 1\}$ such that (i) nodes not colored $\perp$ induce a properly colored subgraph and (ii) nodes colored $\perp$ induce connected components whose size is bounded above by $c_2 \log^{c_3} n$ with probability at least $1 - 1/n$.*

Proof: With respect to the execution of iterations of RANDCOLSTEP, the *state* of a node u is its current color palette C_u and its current color choice $c(u)$. If $c(u) = \perp$, then u has not colored itself; otherwise, $c(u)$ is a permanently assigned color that node u has given itself. To figure out the state of node u after T iterations of RANDCOLSTEP, it suffices to know (i) the state of u and its neighbors after $T - 1$ iterations of RANDCOLSTEP and (ii) at most $\lceil \log \Delta \rceil$ random bits associated with each of these nodes so that their random color choices in iteration T can be determined. Stated differently, it suffices to know (i) the subgraph $G_L[B(u, 1)]$ and (ii) the state of each node in $B(u, 1)$ after $T - 1$ iterations of RANDCOLSTEP. This in turn can be computed from (i) the subgraph $G_L[B(u, 2)]$ and (ii) the state of all nodes in $B(u, 2)$ after $T - 2$ iterations of RANDCOLSTEP. Continuing inductively, we conclude that in order to know the state of node u after T iterations of RANDCOLSTEP, it suffices to know $G_L[B(u, T)]$, where each node v in $B(u, T)$ is labeled with an $(\mathtt{ID}_v, \mathtt{RS}_v)$-pair, where $\mathtt{RS}_v$ is a random bit string of length $T \cdot \lceil \log \Delta \rceil$. □

Now we focus on the problem of each node gathering $G_L[B(u, T)]$ and show that this problem can be solved in $O(\log \log \log n)$ rounds, given that $T = O(\log \log n)$ and $\Delta < \log^4 n$.

Lemma 9. *There is a Congested Clique algorithm running on an n-node input graph G with maximum degree $\Delta < \log^4 n$ that terminates in $O(\log\log\log n)$ rounds at the end of which, every node u knows $G_L[B(u,T)]$.*

Proof: The algorithm starts with each node u broadcasting its degree in G to all nodes in V. This enables every node to locally compute Δ and also a random bit string $\mathtt{RS}_u$ of length $T \cdot \lceil \log \Delta \rceil$. After computing $\mathtt{RS}_u$, each node u sends to each neighbor in G the pair $(\mathtt{ID}_u, \mathtt{RS}_u)$. Now each node u is in possession of the collection of $(\mathtt{ID}_v, \mathtt{RS}_v)$-pairs for all neighbors v. Each node u now has a goal of sending this collection to every neighbor. Note that the total volume of information that u wishes to send out is bounded above by Δ^2 (measured in $O(\log n)$-sized words). Also, each node u is the destination for at most Δ^2 words. Since $\Delta^2 = o(n)$, using Lenzen's routing protocol [12], each node can successfully send its entire collection of $(\mathtt{ID}, \mathtt{RS})$-pairs to all neighbors in constant rounds. Based on this received information, each node u can construct $G_L[B(u,1)]$.

Proceeding inductively, suppose that each node u has gathered $G_L[B(u,t)]$, where $1 \leq t < T$. We now show that in an additional constant rounds, u can gather $G_L[B(u,2t)]$. First note that $|B(u,t)| \leq \Delta^{t+1}$ for any node $u \in V$. Therefore, $G_L[B(u,t)]$ can be completely described using $O(\Delta^{t+2})$ words of information. In order to compute $G_L[B(u,2t)]$, each node u sends $G_L[B(u,t)]$ to each node in $B(u,t)$. A node u, on receiving $G_L[B(v,t)]$ for all nodes v in $B(u,t)$, can perform a local computation to determine $G_L[B(u,2t)]$. Note that the total volume of information that u needs to send out during this communication is $O(\Delta^{2t+3})$ words. By a symmetric reasoning, each node u is the destination for at most $O(\Delta^{2t+3})$ words of information. Since $\Delta < \log^4 n$ and $t < T = O(\log\log n)$, $\Delta^{2t+3} = o(n)$ and therefore using Lenzen's routing protocol, each node u can send $G_L[B(u,t)]$ to each node in $B(u,t)$ in constant rounds.

Since the goal of the algorithm is for each node u to learn $G_L[B(u,T)]$, where $T = O(\log\log n)$, it takes $O(\log\log\log n)$ iterations of the above described inductive procedure to reach this goal. The result follows from the fact that each iteration involves a constant number of communication rounds. □

An immediate consequence of Lemmas 8 and 9 is that there is a Congested Clique algorithm running on an n-node input graph G with maximum degree $\Delta < \log^4 n$ that terminates in $O(\log\log\log n)$ rounds at the end of which, every node u has assigned itself a color $c(u) \in \{\perp\} \cup \{1, 2, \ldots, \Delta+1\}$ such that (i) nodes not colored $\perp$ induce a properly colored subgraph and (ii) nodes colored $\perp$ induce connected components whose size is bounded above by $O(\text{poly}(\log n))$ with probability at least $1 - 1/n$. This brings us to Stage 2 of our algorithm. The first task in this stage is to distribute information about uncolored nodes (i.e., nodes u with $c(u) = \perp$) such that each connected component in the subgraph induced by uncolored nodes ends up at a node in the network. To perform this task in $O(\log\log\log n)$ rounds, we construct a complete, edge-weighted graph in which an edge $\{u,v\}$ has weight $w(u,v) = 1$ if $\{u,v\} \in E$ and $c(u) = c(v) = \perp$ and has weight n otherwise. Thus, edges in the subgraph of G induced by uncolored nodes have weight 1 and edges connecting all other pairs of nodes have weight n. This complete, edge-weighted graph serves as an input to the MST algorithm of Lotker et al. Note that this input is distributed across the network with each node having knowledge of the weights of all $n-1$ edges incident on it. Also note that this knowledge can be acquired by all nodes after just one round of communication. As mentioned earlier, the Lotker et al. MST algorithm runs in $O(\log\log n)$ rounds. Since we are not interested in computing an MST, but only in identifying connected components, we do not have to run the Lotker et al. algorithm to completion.

The Lotker et al. algorithm runs in phases, taking constant number of communication rounds per phase. At the end of phase $k \geq 0$, the algorithm has computed a partition $\mathcal{F}^k = \{F_1^k, F_2^k, \ldots, F_{m_k}^k\}$ of the nodes of G into *clusters*. Furthermore, for each cluster $F \in \mathcal{F}^k$, the algorithm has computed a spanning tree $T(F)$. The correctness of the algorithm is ensured by the fact that each tree $T(F)$ is a subgraph of the MST. It is worth noting that every node in the network knows the partition $\mathcal{F}^k$ and the collection $\{T(F) \mid F \in \mathcal{F}^k\}$ of trees. Suppose that the minimum size cluster in $\mathcal{F}^k$ has size N. The $O(\log\log n)$ running time of the Lotker et al. algorithm arises from the fact that in each phase the algorithm merges clusters and at the end of Phase $k+1$ the smallest cluster in $\mathcal{F}^{k+1}$ has size at least N^2. Thus the size of the smallest cluster "squares" in each phase and therefore it takes $O(\log\log n)$ rounds to get to the stage where the smallest cluster has size n, at which point there is only one cluster F and $T(F)$ is the MST.

We are interested in executing T phases of the Lotker et al. algorithm so that the size of the smallest cluster in $\mathcal{F}^T$ is at least the size of the largest connected component induced by uncolored nodes. Since the size of the largest connected component in the graph induced by uncolored nodes is $O(\text{poly}(\log n))$, it takes only $T = O(\log\log\log n)$ phases to reach such a stage. Let $\mathcal{F}^T = \{F_1^T, F_2^T, \ldots, F_m^T\}$ be the partition of the nodes of G into clusters at the end of T phases of the Lotker et al. algorithm.

Lemma 10. *Let C be a connected component in the subgraph induced by uncolored nodes. Then $C \subseteq F_i^T$ for some i.*

Proof: To obtain a contradiction suppose that $C \cap F_i^T \neq \emptyset$ and $C \cap F_j^T \neq \emptyset$ for some $1 \leq i \neq j \leq m$. Then there is an edge of weight 1 connecting a node in F_i^T and a node in F_j^T. Since $|F_i^T| \geq |C|$, the tree $T(F_i^T)$ contains an edge

of weight n. Thus at some point in the Lotker et al. algorithm, it chose to merge clusters using an edge of weight n when it could have used an edge of weight 1. This contradicts the behavior of the Lotker et al. algorithm. □

The rest of Stage 2 is straightforward. One node, say u^*, considers each $F \in \mathcal{F}^T$ and deletes all edges of weight n from $T(F)$. This will result in F splitting up into smaller clusters; these clusters are the connected components of the subgraph of G induced by uncolored nodes. Note that at this point we think of a connected component as simply a subset of nodes. Node u^* then ships off each connected component to a distinct node, possibly the node with the smallest ID in that component. This takes constant number of rounds via the use of Lenzen's routing protocol. Suppose that a node u has received a connected component C. Node u then contacts the nodes in C to find out (i) all edges connecting pairs of nodes in C, and (ii) the current palettes C_v for each node $v \in C$. Since $|C|$ is polylogarithmic in size and $\Delta < \log^4 n$, it is easy to see that all of this information requires polylogarithmic number of bits to represent and therefore can be communicated to u in constant number of rounds via Lenzen's routing protocol. Node u then colors each node $v \in C$ using a color from its palette C_v such that the graph induced by C is properly colored. This completes Stage 2 and we have a $(\Delta+1)$-coloring of G.

Lemma 11. *Given an n-node graph G with maximum degree $\Delta \leq \log^4 n$, Algorithm* LOWDEGCOL *computes a proper $(\Delta+1)$-coloring in $O(\log\log\log n)$ rounds in the Congested Clique model.*

Combining Lemmas 6 and 7 along with Lemma 11 gives the following theorem.

Theorem 1. *Given an n-vertex input graph $G = (V, E)$ with maximum degree $\Delta \geq \log^4 n$, Algorithm* HIGHDEGCOL *computes an $O(\Delta)$-coloring in $O(1)$ rounds (in expectation) in the Congested Clique model. For arbitrary Δ, an $O(\Delta)$-coloring can be computed in $O(\log\log\log n)$ rounds in expectation in the Congested Clique model.*

3 MapReduce Algorithms from Congested Clique Algorithms

In this section, we prove a *simulation* theorem establishing that Congested Clique algorithms (with fairly weak restrictions) can be efficiently implemented in the MapReduce model. The simulation ensures that a Congested Clique algorithm running in T rounds can be implemented in $O(T)$ rounds (more precisely, $3 \cdot T + O(1)$ rounds) in the MapReduce model, if certain communication and "memory" conditions are met. The technical details of this simulation are conceptually straightforward, but the details are a bit intricate.

We will now precisely define restrictions that we need to place on Congested Clique algorithms in order for the simulation theorem to go through. We assume that each node in the Congested Clique possesses a word-addressable memory whose words are indexed by the natural numbers. For an algorithm $\mathcal{A}_{CC}$ running in the Congested Clique, let $I_u^{(j)} \subset \mathbb{N}$ be the set of memory addresses *used* by node u during the local computation in round j (not including the sending and receipt of messages).

After local computation in each round, each node in the Congested Clique may send (or not send) a distinct message of size $O(\log n)$ to each other node in the network. In defining notation, we make a special distinction for the case where a node u sends in the *same* message to every other node v in a particular round; i.e., node u sends a *broadcast* message. The reason for this distinction is that broadcasts can be handled more efficiently on the receiving end in the MapReduce framework than can distinct messages sent by u. Let $m_{u,v}^{(j)}$ denote a message sent by node u to node v in round j and let $D_u^{(j)} \subseteq V$ be the set of destinations of messages sent by node u in round j. Let $M_u^{(j)} = \{m_{u,v}^{(j)} \colon v \in D_u^{(j)} \subset V\}$ be the set of messages *sent* by node u in round j of algorithm $\mathcal{A}_{CC}$, except let $M_u^{(j)} = \emptyset$ if u has chosen to broadcast a message $b_u^{(j)}$ in round j. Similarly, let $\overline{M}_u^{(j)} = \{m_{v,u}^{(j)} \colon u \in D_v^{(j)}$ and v is not broadcasting in round $j\}$ be the set of messages *received* by node u in round j, except that we exclude messages $b_v^{(j)}$ from nodes v that have chosen to broadcast in round j. We say that $\mathcal{A}_{CC}$, running on an n-node Congested Clique, is (K, N)-*lightweight* if

(i) for each round j (in the Congested Clique), $\sum_{u \in V}(|\overline{M}_u^{(j)}| + |I_u^{(j)}|) = O(K)$;
(ii) there exists a constant C such that for each round j and for each node u, $I_u^{(j)} \subseteq \{1, 2, \ldots, \lceil C \cdot N \rceil\}$; and
(iii) each node u performs only polynomial-time local computation in each round.

In plain language: no node uses more than $O(N)$ memory for local computation during a round; the total amount of memory that all nodes use and the total volume of messages nodes receive in any round is bounded by $O(K)$. Regarding condition (iii), traditional models of distributed computation such as the $\mathcal{CONGEST}$ and $\mathcal{LOCAL}$ models allow nodes to perform arbitrary local computation (e.g., taking exponential time), but since the MapReduce model requires mappers and reducers to run in polynomial time, we need this extra restriction.

Theorem 2. *Let ϵ, c satisfy $0 \leq \epsilon \leq c$, and let $G = (V, E)$ be a graph on n vertices having $O(n^{1+c})$ edges. If $\mathcal{A}_{CC}$ is a $(n^{1+c}, n^{1+\epsilon})$-lightweight Congested Clique-model algorithm running on input G in T rounds, then $\mathcal{A}_{CC}$ can be implemented in the MapReduce model with $n_r = n^{c-\epsilon}$ machines and $m_r = \Theta(n^{1+\epsilon})$ (words of) memory per machine such that the implementation runs in $O(T)$ Map-Shuffle-Reduce rounds on G.*

Proof: The simulation that will prove the above theorem contains two stages: the *Initialization* stage and the *Simulation* stage. In the Initialization stage, the input to the MapReduce system is transformed from the assumed format (an unordered list of edges and vertices of G) into a format in which each piece of information, be it an edge, node, or something else, that is associated with a node of G is gathered at a single machine. After this gathering of associated information has been completed, the MapReduce system can emulate the execution of the Congested Clique algorithm.
Initialization stage. Input (in this case, the graph G) in the MapReduce model is assumed to be presented as an unordered sequence of tuples of the form $(\varnothing, u)$, where u is a vertex of G, or $(\varnothing, (u, v))$, where (u, v) is an edge of G. The goal of the Initialization stage is to partition the input G among the n_r reducers such that each reducer r receives a subset $P_r \subseteq V$ and all edges E_r incident on nodes in P_r such that $|P_r| + |E_r|$ is bounded above by $O(n^{1+\epsilon})$. This stage can be seen as consisting of two tasks: (i) every reducer r learns the degree $\deg_G(u)$ of every node u in G and (ii) every reducer computes a partition (the same one) given by the partition function $F_0 : V \longrightarrow \{1, 2, \ldots, n_r\}$, defined by

$$F_0(x) = \begin{cases} 1, & \text{if } x = 1 \\ F_0(x-1), & \text{if } \sum_{v \in L(x)} \deg_G(v) \leq n^{1+\epsilon}, \\ F_0(x-1) + 1, & \text{otherwise} \end{cases}$$

Here $L(x) = \{j < x : F_0(j) = F_0(x-1)\}$. All nodes in the same group in the partition are mapped to the same value by F_0 and will be assigned to a single reducer. Since the degree of each node is bounded above by n, it is easy to see that for any $r \in \{1, 2, \ldots, n_r\}$, $F_0^{-1}(r)$ is a subset of nodes of G such that $|F_0^{-1}(r)| + \sum_{u \in F_0^{-1}(r)} \deg_G(u)$ is $O(n^{1+\epsilon})$. Each of the two tasks mentioned above can be implemented in a (small) constant number of MapReduce rounds as follows.

- **Map 1:** In Map phase 1, for each tuple $(\varnothing, u)$, a mapper chooses a random reducer r and emits the tuple (r, u). For each tuple $(\varnothing, (u, v))$, a mapper again chooses a random reducer r and emits the tuple $(r, (u, v))$. Because the reduce keys are chosen at random, with high probability (actually, exponentially high probability) each reducer in Reduce phase 1 will receive $O(n^{1+\epsilon})$ tuples.
- **Reduce 1:** In Reduce phase 1, a reducer r receives tuples whose values consist of some collection $P_r \subseteq V$ of vertices and some collection $E_r \subseteq E$ of edges of G. For each value consisting of a vertex u, a reducer r re-emits the tuple (r, u), and for each value consisting of an edge (u, v), reducer r re-emits the tuple$(r, (u, v))$. In addition, a reducer r emits, for each vertex u such that reducer r received an edge (u, v) or (v, u), a tuple $(r, u, d_{r,u})$, where $d_{r,u}$ is total number of edges received by reducer r containing u. (In other words, $d_{r,u}$ is the partial degree of u seen by reducer r.)
- **Map 2:** In Map phase 2, mappers again load-balance tuples containing vertices or edges as values across the reducers uniformly at random (an action which is successful w.h.p.), as in Map phase 1. In addition, when a mapper processes a tuple of the form $(r, u, d_{r,u})$, it emits the tuple $((u \bmod n_r), u, d_{r,u})$. Here $u \bmod n_r$ refers to the reduction of the identifier of node u modulo the number of reducers, n_r. There are at most $n \cdot n_r = O(n^{1+c-\epsilon})$ such tuples, and thus (i) each reducer is the destination of $O(n)$ such tuples (of the form $((u \bmod n_r), u, d_{r,u})$); and (ii) all tuples containing a partial degree sum of node u among their values are given the same key and thus sent to the same reducer during the second MapReduce round.
- **Reduce 2:** In Reduce phase 2, a reducer r again re-emits tuples (r, u) and $(r, (u, v))$ for each vertex or edge received as a value. For tuples of the form $(r, u, d_{r',u})$, reducer r aggregates the partial degree sums of u to compute the full degree $\deg_G(u)$ of u in G, and emits the tuple $(r, u, \deg_G(u))$.
- **Map 3:** In Map phase 3, mappers once again load-balance tuples containing vertices or edges as values across the reducers as in Map phases 1 and 2. For each tuples of the form $(r, u, \deg_G(u))$, a mapper emits n_r tuples $(r_1, u, \deg_G(u))$, $(r_1, u, \deg_G(u))$, $\ldots$, $(r_{n_r}, u, \deg_G(u))$ – one for each reducer. Thus, for each reducer, exactly n tuples containing (full) degree information are emitted – one for each vertex of G.
- **Reduce 3:** In Reduce phase 3, a reducer r now has access to the degrees of all vertices of G and can thus compute the partition function F_0 defined earlier. Then, for each node u received, a reducer r outputs the tuple $(r, F_0(u), u)$, and for each edge (u, v) received, a reducer r outputs the tuples $(r, F_0(u), (u, v))$ and $(r, F_0(v), (u, v))$.
- In addition to "packaging" the vertex and edge information of G so that incident edges of a node u can be collected at the reducer $F_0(u)$ assigned to simulate computation at u, reducers must also emit tuples which allow both (i)

the currently collected degrees of each vertex in G and (ii) the partition function F_0 to be propagated forward through the rounds of the MapReduce computation. Fortunately this is straightforward: for each degree tuple $(r, u, \deg_G(u))$ received by reducer r, reducer r re-emits the same tuple. As well, $F_0 : V \longrightarrow \{1, \ldots, n_r\}$ can be fully described by n pairs $(v, F_0(v))$, and so reducer r emits the n tuples $(r, v, F_0(v))$, which will allow reducer r to "remember" the partition function $F_0(\cdot)$ in the next round. Observe that the totality of the memory required to support knowledge of the partition function and all degrees in G is $O(n)$, and thus fits into the memory of a reducer without any trouble.

– **Map 4:** Finally, in Map phase 4, a mapper receives and processes two different tuple formats: (i) tuples of the form (r, r', z), where r' is another reducer index and z is some information (of length $O(1)$ words) representing either a vertex or an edge; and (ii) tuples of the form (r, v, z), where v is a vertex identifier and z is either a degree value or a reducer identifier. In case (i) (tuples of the form (r, r', z)), a mapper emits the tuple (r', z). In case (ii) (tuples of the form (r, v, z), a mapper simply outputs the same tuple (r, v, z) unchanged.
– After the Map phase of the round 4 of the MapReduce computation has completed, the Initialization phase is complete, and the simulation of $\mathcal{A}_{CC}$ is ready to begin.

Simulation stage. At a high level, a Reduce phase serves as the "local computation" phase of the Congested Clique simulation, whereas a Map phase (together with the subsequent shuffle phase) serves as the "communication" phase of the simulation. However, there is, in general, a constant-factor slow-down because it may be that the sending and receiving of messages in $\mathcal{A}_{CC}$ could cause the subset of nodes assigned to a reducer to aggregate more than $O(n^{1+\epsilon})$ memory, necessitating a re-partitioning of the nodes among the reducers so as not to violate the memory-per-machine constraint.

Recall that $I_u^{(i)}$ denotes the set of memory addresses used by a node u in round i of $\mathcal{A}_{CC}$. Let $h_{u,j}^{(i)}$ be the value of word $j \in I_u^{(i)}$ in the memory of node u after node u has completed local computation in round i of $\mathcal{A}_{CC}$, but before messages have been sent and received in this round. For $i > 0$, define a tuple set

$$\mathcal{H}_u^{(i)} = \{(F_i(u), (u, i, h_{u,j}^{(i)})) \colon j \in I_u^{(i)}\}$$

where $F_i(\cdot)$ is the partition function used in round i. Like F_0, defined in the Initialization stage, F_i partitions G into n_r groups, one per reducer, so that reducer memory constraints are not violated in round i. The collection of tuples $\mathcal{H}_u^{(i-1)}$ is a representation, in the MapReduce key-value format, of the information necessary to simulate the computations of node u in round i of the Congested Clique algorithm $\mathcal{A}_{CC}$. The use of $F_i(u)$ as the key in each of the tuples in $\mathcal{H}_u^{(i)}$ ensures that all information needed to simulate a local computation at u in $\mathcal{A}_{CC}$ goes to the same reducer. Additionally, note that the inclusion of the identifier of u with the values allows the words from u's memory to be reassembled and distinguished from information associated with other nodes $v \in F_i^{-1}(u)$. We assume that $\mathcal{H}_u^{(0)}$ is the information in tuple format that node u has initially about graph G. In other words, $\mathcal{H}_u^{(0)} = \{(F_0(u), u)\} \cup \{(F_0(u), (u, v)) : v \text{ is a neighbor of } u\}$.

Once an initial partition function $F_0(\cdot)$ has been computed and the initial collections $\mathcal{H}_u^{(0)}$ have been assembled the main goals of our simulation algorithm are to (i) provide a mechanism for transforming $\mathcal{H}_u^{(i-1)}$ into $\mathcal{H}_u^{(i)}$ during the reduce phase of a MapReduce round; and (ii) provide a means of transmitting messages to reducers of a subsequent round (corresponding to messages transmitted in the Congested Clique at the end of each round). Since we assume messages to be sent and received after local computation has occurred during a Congested Clique round, $\mathcal{M}_u^{(i)}$ can be determined from $\mathcal{H}_u^{(i)}$; in turn, $\mathcal{H}_u^{(i)}$ is a function of $\mathcal{H}_u^{(i-1)}$ and $\overline{\mathcal{M}}_u^{(i-1)}$.

We describe the details of the simulation of a single round (round i) of a Congested Clique algorithm $\mathcal{A}_{CC}$ below. Let $j = 3i - 1$. Round i of $\mathcal{A}_{CC}$ is simulated by three MapReduce rounds (a total of six Map or Reduce phases) – Reduce $j-1$, Map j, Reduce j, Map $j+1$, Reduce $j+1$, and Map $j+2$. We assume inductively that as input to Reduce phase $j-1$ below, each reducer receives, in addition to data tuples, $O(n)$ metadata tuples containing a description of a partition function $F_{i-1}(\cdot)$ such that for each r, $\sum_{u \in P_r}(|\mathcal{H}_u^{(i-1)}| + |\overline{\mathcal{M}}_u^{(i-1)}|) = O(n^{1+\epsilon})$, where $P_r = F_{i-1}^{-1}(r)$.

– **Reduce phase** $j-1$**:** In Reduce phase $j-1$, a reducer r receives input consisting of $\mathcal{H}_u^{(i-1)}$ together with $\overline{\mathcal{M}}_u^{(i-1)}$ for each $u \in P_r$; for each such u, reducer r performs the following steps:
  (i) Reducer r simulates the local computation of Round i of $\mathcal{A}_{CC}$ at u.
  (ii) Reducer r computes $\mathcal{H}_u^{(i)}$ from $\mathcal{H}_u^{(i-1)}$ and $\overline{\mathcal{M}}_u^{(i-1)}$, *but does not yet output* any tuples of $\mathcal{H}_u^{(i)}$; rather, reducer r outputs only a tuple (r, u, s_u) containing the size of the information $s_u = |\mathcal{H}_u^{(i)}|$.

(iii) Reducer r computes $\mathcal{M}_u^{(i)}$ from $\mathcal{H}_u^{(i)}$, but again, *does not output* any elements of $\mathcal{M}_u^{(i)}$. Reducer r then computes, for each $v \in V$, the aggregate count $c_{r,v}$ of messages emanating from nodes in P_r and destined for v, and outputs the tuple $(r, v, c_{r,v})$.

(iv) Reducer r outputs the exact same tuples it received as input, $\mathcal{H}_u^{(i-1)}$ and $\overline{\mathcal{M}}_u^{(i-1)}$.

– **Map phase** j**:** Before message tuples can be generated and aggregated (as a collection $\overline{M}_u^{(i)}$ at reducer $F(u)$) a rebalancing of the nodes to reducers must be performed to ensure that the reducer-memory constraint is not violated. In Map phase j, a mapper forwards tuples from either a $\mathcal{H}_u^{(i-1)}$ or a $\overline{\mathcal{M}}_u^{(i-1)}$ through unchanged. However, for each tuple of the form $(r, u, c_{r,u})$, a mapper outputs the tuple $(u \bmod n_r, u, c_{r,u})$. In addition, for each tuple of the form (r, u, s_u), a mapper outputs n_r tuples (r', u, s_u) – one for each reducer r' – so that every reducer can know the future size of $\mathcal{H}_u^{(i)}$.

– **Reduce phase** j**:** In Reduce phase j, a reducer r receives as input nearly the exact same input (and output) of reducer r in the previous MapReduce round – the union of $\mathcal{H}_u^{(i-1)}$ and $\overline{\mathcal{M}}_u^{(i-1)}$ for each $u \in P_r$ – except that instead of receiving tuples of the form $(r, u, c_{r,u})$ for each $u \in V$, reducer r receives *all* partial message counts for the subset of vertices u for which $u \bmod n_r = r$; as well, each reducer receives n tuples of the form (r, u, s_u) describing the amount of memory required by node u in round i of $\mathcal{A}_{CC}$. Reducer r aggregates tuples of the form $(u \bmod n_r, u, c_{r,u})$ and outputs $(r, u, |\overline{M}_u^{(i)}|)$, since $|\overline{M}_u^{(i)}|$ is precisely the sum of the partial message counts $c_{r,u}$. (Notice that a reducer r receives $O(n)$ such tuples.) Reducer r forwards all other tuples through unchanged to the next MapReduce round.

– **Map phase** $j+1$**:** In Map phase $j+1$, a mapper continues to forward all tuples through unchanged to Reduce phase $j+1$, except that for each tuple of the form $(r, u, |\overline{M}_u^{(i)}|)$, a mapper outputs n_r tuples $(r', u, |\overline{M}_u^{(i)}|)$ – one for each reducer r'. In this way, each reducer in Reducer phase $j+1$ can come to know all n message counts for each node $u \in V$.

– **Reduce phase** $j+1$**:** In Reduce phase $j+1$, each reducer receives all n message counts (for each node $u \in V$) in addition to the sizes s_u of the state needed by each node u in round i of $\mathcal{A}_{CC}$. Each reducer thus has enough information to determine the next partition function $F_i : V \longrightarrow \{1, \ldots, n_r\}$, defined by

$$F_i(x) = \begin{cases} 1, & \text{if } x = 1 \\ F_i(x-1), & \text{if } \sum_{v \in L(x)} (s_v + |\overline{M}_v^{(i)}|) \leq n^{1+\epsilon}, \\ F_i(x-1) + 1, & \text{otherwise} \end{cases}$$

Here $L(x) = \{v \mid v < x \text{ and } F_i(v) = F_i(x-1)\}$. After determination of the new partition function F_i, reducers are now able to successfully output the "packaged memory" $\mathcal{H}_u^{(i)}$ of round i of $\mathcal{A}_{CC}$, as well as the new messages $m_{u,v}^{(i)}$ sent in round i, because the new partition function F_i is specifically designed to correctly load-balance these tuple sets across the reducers while satisfying the memory constraint. Therefore:

(i) Reducer r now simulates the local computation at each $u \in P_r$ and thus outputs the set $\mathcal{H}_u^{(i)}$ (which can be computed from $\mathcal{H}_u^{(i-1)}$ and $\overline{\mathcal{M}}_u^{(i-1)}$). It is important to recall here that because mappers operate on key-value pairs one at a time in the MapReduce model, there is no restriction on the size of the output from any reducer r in any MapReduce round (other than that it be polynomial). [9] Therefore, a reducer r may output (and thus free-up its memory) each tuple set $\mathcal{H}_u^{(i)}$ as it is created (as reducer r processes the nodes in P_r one at a time), and so there is no concern about reducer r attempting to maintain in memory all sets $\mathcal{H}_u^{(i)}$ for $u \in P_r$ at once. Note that $\mathcal{H}_u^{(i)}$, as generated by a reducer r, should contain tuples of the form $(r, F_i(u), u, h_{u,l}^{(i)})$ so that mappers in MapReduce round $j+2$ can correctly deliver $\mathcal{H}_u^{(i)}$ to reducer $F_i(u)$. Recall that $h_{u,l}^{(i)}$ denotes the contents of the word with address l in node u's memory at the end of local computation in round i.

(ii) As a reducer r processes, and simulates the computation at, each node $u \in P_r$ one at a time, generating $\mathcal{H}_u^{(i)}$, reducer r also uses $\mathcal{H}_u^{(i)}$ to generate the messages $M_u^{(i)}$ to be sent by node u in round i of $\mathcal{A}_{CC}$. Reducer r encapsulates $M_u^{(i)}$ in the tuple set $\mathcal{M}_u^{(i)}$ and outputs it alongside $\mathcal{H}_u^{(i)}$ before moving on to the next node in P_r. As with $\mathcal{H}_u^{(i)}$, tuples in $\mathcal{M}_u^{(i)}$ should initially be generated by a reducer r in the form $(r, F_i(v), u, v, m_{u,v}^{(i)})$ so that mappers in MapReduce round $j+2$ can correctly deliver the set $\overline{\mathcal{M}}_v^{(i)}$ to reducer $F_i(v)$.

(iii) Lastly regarding the simulation procedure, whenever a node $u \in P_r$ being simulated broadcasts a message $b_u^{(i)}$, reducer r outputs the tuple $(r, u, b_u^{(i)})$.

(iv) After simulation of each node $u \in P_r$ is complete, reducer r also outputs a description of the new partition function F_i.

– **Map** $j+2$**:** In Map phase $j+2$, a mapper simply transforms the key in a data tuple as appropriate: for each tuple $(r, F_i(u), u, h_{u,l}^{(i)})$, a mapper simply emits the tuple $(F_i(u), u, h_{u,l}^{(i)})$; for each tuple $(r, F_i(v), u, v, m_{u,v}^{(i)})$, a mapper simply emits the tuple $(F_i(v), u, v, m_{u,v}^{(i)})$. The exception to this is that tuples $(r, u, b_u^{(i)})$ containing broadcast messages are expanded: for each, a mapper emits n_r tuples $(r', u, b_u^{(i)})$ – one for each reducer r' – so that every reducer in Reducer phase $j+2$ receives a single copy of each message broadcast during round i of $\mathcal{A}_{CC}$.
– Tuples carrying metadata describing the (new) partition function F_i are forwarded unchanged, because there already exists one copy of each such metadata tuple for each reducer, and there need be only one such copy per reducer as well. After Map phase $j+2$, tuples from the sets $\mathcal{H}_u^{(i)}$ and $\overline{\mathcal{M}}_u^{(i)}$ have been emitted with keys $F_i(u)$, and for each broadcast message $b_u^{(i)}$, one tuple containing a copy of $b_u^{(i)}$ has been emitted for each reducer as well; thus, in Reduce phase $j+2$, simulation of round $i+1$ of algorithm $\mathcal{A}_{CC}$ can begin.

It remains to comment on the memory-per-machine constraint which must be satisfied during each MapReduce round. Observe that, inductively, for each r, the sum $\sum_{u \in P_r}(|\mathcal{H}_u^{(i-1)}| + |\overline{\mathcal{M}}_u^{(i-1)}|) = O(n^{1+\epsilon})$. These data tuples are forwarded unchanged until Reduce phase $j+1$, in which the new partition function $F_i(\cdot)$ for the next round of simulation is computed, and then collectively $\mathcal{H}_u^{(i-1)}$ and $\overline{\mathcal{M}}_u^{(i-1)}$ are transformed into $\mathcal{H}_u^{(i)}$ and $\mathcal{M}_u^{(i)}$. By construction of the partition functions F_{i-1} and F_i, and by the assumption that $\mathcal{A}_{CC}$ is a $(n^{1+c}, n^{1+\epsilon})$-*lightweight* algorithm, it follows that these data tuples are never present on any reducer a number that exceeds $\Theta(n^{1+\epsilon})$. Secondly, it should be mentioned that because broadcast messages are not duplicated at any reducer r, no reducer will ever receive more than $n = O(n^{1+\epsilon})$ tuples containing broadcast messages. Thirdly, tuples containing state or message counts are never present in a number exceeding n at any reducer, and *partial* message counts are explicitly load-balanced so that only $O(n)$ such information is passed to a single reducer as well. Finally, metadata tuples describing a partition function never exceed $\Theta(n)$ on any reducer because the domain of each partition function has size n. □

4 Coloring in the MapReduce Framework

Using the simulation theorem of Section 3, we can simulate Algorithm HIGHDEGCOL in the MapReduce model and thereby achieve an $O(\Delta)$-coloring MapReduce algorithm running in expected-$O(1)$ rounds. As in Lattanzi et al. [11], we consider graphs with $\Omega(n^{1+c})$ edges, $c > 0$.

Theorem 3. *When the input graph G has $\Omega(n^{1+c})$ edges, and $0 \leq \epsilon < c$, there exists an $O(\Delta)$-coloring algorithm running in the MapReduce model with $\Theta(n^{c-\epsilon})$ machines and $\Theta(n^{1+\epsilon})$ memory per machine, and having an expected running time of $O(1)$ rounds.*

Proof: It is easy to examine the lines of code in Algorithm HIGHDEGCOL to ascertain that the total amount of non-broadcast communication in any round in bounded above by $O(n^{1+c})$. Specifically, the total non-broadcast communication corresponding to only two lines of code – Lines 6 and 11 – can be as high as $\Theta(n^{1+c})$. For all other lines of code, the volume of total non-broadcast communication is bounded by $O(n)$. Similarly, it is easy to examine the lines of code in Algorithm HIGHDEGCOL to verify that the total memory (in words) used by all nodes for their local computations in any one round is bounded above by $O(n^{1+c})$. Finally, it is also easy to verify that the maximum amount of memory used by a node in any round of computation is $O(n)$.

Thus, Algorithm HIGHDEGCOL is an (n^{1+c}, n)-lightweight algorithm on a Congested Clique and applying the Simulation Theorem (Theorem 1) to this algorithm yields the claimed result. □

It is worth emphasizing that the result holds even when $\epsilon = 0$; in other words, even when the per machine memory is $O(n)$, the algorithm can compute an $O(\Delta)$-coloring in $O(1)$ rounds. This is in contrast with the results in Lattanzi et al. [11], where $O(1)$-round algorithms were obtained (e.g., for maximal matching) with $n^{1+\epsilon}$ per machine memory, only when $\epsilon > 0$. In their work, setting $\epsilon = 0$ (i.e., using $\Theta(n)$ memory per machine) resulted in $O(\log n)$ round algorithms.

We end with the following corollary that is an immediate consequence of Theorem 3.

Corollary 1. *The problem of computing an $O(\Delta)$-coloring for an n-node graph with maximum degree Δ and at least $\Omega(n^{1+c})$ edges, for $c > 0$ is in $\mathcal{MRC}^0$.*

5 Conclusions

The results in this paper connect two models that are usually studied by different research communities. In general, it would be interesting to see if this connection has benefits beyond those discussed in the paper. Also, it would be be interesting to study differences between these two models. For example, the Congested Clique model allows nodes to remember arbitrary amount of information from one round to the next. Does this give the Congested Clique model a provable advantage over the "stateless" MapReduce model?

For the "small Δ" case, i.e., when $\Delta = O(\text{poly}(\log n))$, our paper presents an $O(\log\log\log n)$-round $(\Delta+1)$-coloring algorithm on a Congested Clique. One question that interests us is whether $O(1)$ rounds will suffice to compute an $O(\Delta)$-coloring even when Δ is small?

Following the lead of Lattanzi et al. [11], we have assumed that each machine in the MapReduce model contains at least $\Omega(n)$ memory for processing an n-node graph. Relaxing this assumption is interesting and leads to the question of whether for some $\epsilon > 0$, $O(1)$ MapReduce rounds would suffice to compute an $O(\Delta)$-coloring, even when the per machine memory is $O(n^{1-\epsilon})$.

References

1. Leonid Barenboim, Michael Elkin, Seth Pettie, and Johannes Schneider. The locality of distributed symmetry breaking. In *Proc. of IEEE FOCS*, 2012.
2. Andrew Berns, James Hegeman, and Sriram V. Pemmaraju. Super-Fast Distributed Algorithms for Metric Facility Location. In Artur Czumaj, Kurt Mehlhorn, Andrew M. Pitts, and Roger Wattenhofer, editors, *ICALP (2)* , volume 7392 of *Lecture Notes in Computer Science*, pages 428–439. Springer, 2012.
3. Andrew Berns, James Hegeman, and Sriram V. Pemmaraju. Super-Fast Distributed Algorithms for Metric Facility Location. *CoRR*, abs/1308.2473, August 2013.
4. Jeffrey Dean and Sanjay Ghemawat. Mapreduce: Simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, January 2008.
5. Danny Dolev, Christoph Lenzen, and Shir Peled. "tri, tri again": Finding triangles and small subgraphs in a distributed setting. In *Proceedings of the 26th International Conference on Distributed Computing*, DISC'12, pages 195–209, Berlin, Heidelberg, 2012. Springer-Verlag.
6. Alina Ene, Sungjin Im, and Benjamin Moseley. Fast clustering using mapreduce. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '11, pages 681–689, New York, NY, USA, 2011. ACM.
7. Joachim Gehweiler, Christiane Lammersen, and Christian Sohler. A Distributed O(1)-approximation Algorithm for the Uniform Facility Location Problem. In *Proceedings of the Eighteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '06, pages 237–243, New York, NY, USA, 2006. ACM.
8. Öjvind Johansson. Simple distributed $(\delta+1)$-coloring of graphs. *Inf. Process. Lett.*, 70(5):229–232, June 1999.
9. Howard Karloff, Siddharth Suri, and Sergei Vassilvitskii. A model of computation for mapreduce. In *Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, pages 938–948, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics.
10. Ravi Kumar, Benjamin Moseley, Sergei Vassilvitskii, and Andrea Vattani. Fast greedy algorithms in mapreduce and streaming. In *Proceedings of the Twenty-fifth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '13, pages 1–10, New York, NY, USA, 2013. ACM.
11. Silvio Lattanzi, Benjamin Moseley, Siddharth Suri, and Sergei Vassilvitskii. Filtering: A method for solving graph problems in mapreduce. In *Proceedings of the Twenty-third Annual ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '11, pages 85–94, New York, NY, USA, 2011. ACM.
12. Christoph Lenzen. Optimal Deterministic Routing and Sorting on the Congested Clique. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing*, PODC '13, pages 42–50, 2013.
13. Zvi Lotker, Boaz Patt-Shamir, and David Peleg. Distributed MST for Constant Diameter Graphs. *Distributed Computing*, 18(6):453–460, 2006.
14. Boaz Patt-Shamir and Marat Teplitsky. The round complexity of distributed sorting: Extended abstract. In *Proceedings of the 30th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, PODC '11, pages 249–256, New York, NY, USA, 2011. ACM.
15. David Peleg. *Distributed Computing: A Locality-Sensitive Approach*, volume 5. Society for Industrial Mathematics, 2000.